

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
10 May 2001 (10.05.2001)

PCT

(10) International Publication Number
WO 01/33774 A1

(51) International Patent Classification⁷: **H04L 12/46, 12/44**, (72) Inventor: TZENG, Shr-Jie; 34263 Mimosa Terrace, Fremont, CA 94555 (US).

(21) International Application Number: **PCT/US00/13397**, (74) Agent: RODDY, Richard, J.; Advanced Micro Devices, Inc., One AMD Place, Mail Stop 68, P.O. Box 3453, Sunnyvale, CA 94088-3453 (US).

(22) International Filing Date: **16 May 2000 (16.05.2000)**

(81) Designated States (*national*): JP, KR.

(25) Filing Language: English

(84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

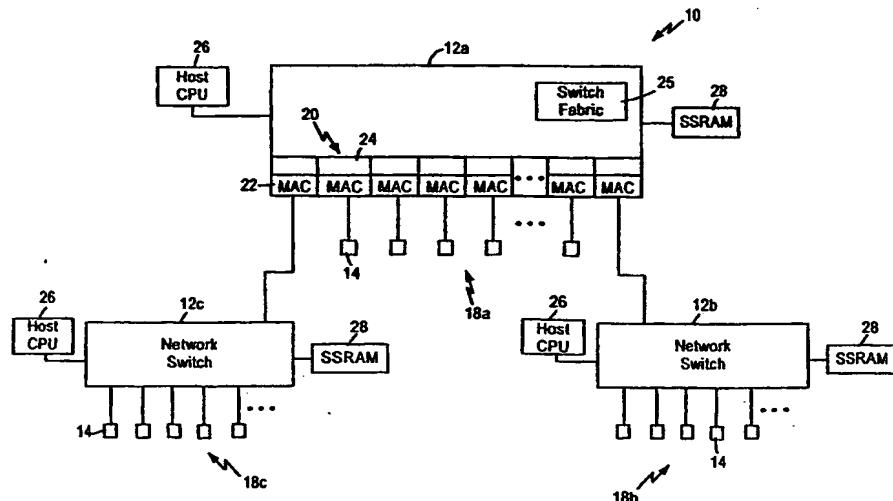
(26) Publication Language: English
(30) Priority Data:
09/430,753 29 October 1999 (29.10.1999) US

Published:
— With international search report.

(71) Applicant: **ADVANCED MICRO DEVICES, INC.**
[US/US]; One AMD Place, Mail Stop 68, P.O. Box 3453,
Sunnyvale, CA 94088-3453 (US).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: APPARATUS AND METHOD FOR IDENTIFYING DATA PACKET TYPES IN REAL TIME ON A NETWORK SWITCH PORT



WO 01/33774 A1

(57) Abstract: A network switch, configured for performing layer 2 and layer 3 switching in an Ethernet (IEEE 802.3) network without blocking of incoming data packets, includes a network switch port having a filter (i.e., a packet classifier module) configured for evaluating an incoming data packet on an instantaneous basis. The filter performs simultaneous comparisons between the incoming data stream of the data packet and multiple templates configured for identifying respective data protocols. Each template is composed of a plurality of min terms, wherein each min term specifies a prescribed comparison operation within a selected data byte of the incoming data packet. The templates may be programmed by a user and stored in an internal min term memory. Moreover, the multiple simultaneous comparisons enable the network switch to perform layer 3 switching for 100 Mbps and gigabit networks without blocking in the network switch.

**APPARATUS AND METHOD FOR IDENTIFYING DATA
PACKET TYPES IN REAL TIME ON A NETWORK
SWITCH PORT**

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

The present invention relates to layer 2 and layer 3 switching of data packets in a non-blocking network switch configured for switching data packets between subnetworks.

BACKGROUND ART

5 Local area networks use a network cable or other media to link stations on the network. Each local area network architecture uses a media access control (MAC) enabling network interface devices at each network node to access the network medium.

10 The Ethernet protocol IEEE 802.3 has evolved to specify a half-duplex media access mechanism and a full-duplex media access mechanism for transmission of data packets. The full-duplex media access mechanism provides a two-way, point-to-point communication link between two network elements, for example between a network node and a switched hub.

15 Switched local area networks are encountering increasing demands for higher speed connectivity, more flexible switching performance, and the ability to accommodate more complex network architectures. For example, commonly-assigned U.S. Patent No. 5,953,335 discloses a network switch configured for switching layer 2 type Ethernet (IEEE 802.3) data packets between different network nodes; a received data packet may include a VLAN (virtual LAN) tagged frame according to IEEE 802.1q protocol that specifies another subnetwork (via a router) or a prescribed group of stations. Since the switching occurs at the layer 2 level, a router is typically necessary to transfer the data packet between subnetworks.

20 Efforts to enhance the switching performance of a network switch to include layer 3 (e.g., Internet protocol) processing may suffer serious drawbacks, as current layer 2 switches preferably are configured for operating in a non-blocking mode, where data packets can be output from the switch at the same rate that the data packets are received. Newer designs are needed to ensure that higher speed switches can provide both layer 2 switching and layer 3 switching capabilities for faster speed 25 networks such as 100 Mbps or gigabit networks.

However, such design requirements risk loss of the non-blocking features of the network switch, as it becomes increasingly difficult for the switching fabric of a network switch to be able to perform layer 3 processing at the wire rates (i.e., the network data rate).

5

SUMMARY OF THE INVENTION

There is a need for an arrangement that enables a network switch to provide layer 2 switching and layer 3 switching capabilities for 100 Mbps and gigabit links without blocking of the data packets.

There is also a need for an arrangement that enables a network switch to provide layer 2 switching and layer 3 switching capabilities with minimal buffering within the network switch that 10 may otherwise affect latency of switched data packets.

There is also a need for an arrangement that enables a network switch to be easily programmable to distinguish between different types of layer 3 data packets so that quality of service (QoS) can be achieved.

There is also a need for an arrangement to enable a network switch port to instantaneously 15 evaluate an incoming data packet and determine a layer 3 or higher protocol, to provide the associated switch fabric with sufficient time to process the incoming data packet according to the detected protocol.

These and other needs are attained by the present invention, where a network switch port includes a filter configured for evaluating an incoming data packet on an instantaneous basis. The 20 filter performs simultaneous comparisons between the incoming data stream of the data packet and multiple templates configured for identifying respective protocols. Each template is composed of a plurality of min terms, where each min term specifies a prescribed comparison operation with a selected data byte of the incoming data packet. The filter accesses the min terms based on the ordering of the data bytes, such that the min terms that are used to compare the first data byte are first 25 accessed for comparison with the first data byte as the first data byte is received; the min terms used to compare the second data byte are then accessed for comparison with the second data byte as the second data byte is received. Hence, the filter simultaneously compares the min terms that correspond to the selected byte of the incoming data packet as the selected byte is received by the network switch port. The results of the comparisons between the min terms and the selected data bytes of the 30 incoming data stream are evaluated by an equation core within the filter, which determines comparison results for the templates and outputs a tag to the switching core, providing the switching core with information on how to process the received data packet. Hence, the switching core has sufficient time to perform layer 2 or layer 3 switching of the received data packet in a non-blocking manner in a high-speed data network.

One aspect of the present invention provides a method of evaluating an incoming data packet at a network switch port. The method includes storing a plurality of templates configured for identifying respective data formats, each template having at least one min term configured for comparing a corresponding prescribed value to a corresponding selected byte of the incoming data packet. The method also includes simultaneously comparing, to the selected byte, the min terms that correspond to the selected byte as the selected byte is received by the network switch port, and generating a comparison result that identifies the incoming data packet, based on the comparisons of the min terms to the data bytes received by the network switch port. The storage of templates configured for identifying respective data formats enables the network switch port to be easily programmable to identify user-defined data formats. Moreover, the simultaneous comparisons of min terms that correspond to a selected byte as the selected byte is received from the network switch port ensures that the incoming data packet can be evaluated in real time, enabling identification of the incoming data packet relative to the user-defined data formats as the data packet is received.

Another aspect of the present invention provides a method of evaluating an incoming data packet at a network switch port. The method includes simultaneously comparing a first byte of the incoming data packet as the first byte is received by the network switch port, with a plurality of templates by comparing the first byte with at least a first min term associated with at least one of the templates, each template configured for identifying a prescribed format within the incoming data packet. The method also includes simultaneously comparing a second byte of the incoming data packet, following the first byte, with the templates by comparing the second byte with at least a second min term associated with at least one of the templates, and generating a comparison result based on min term results from the comparisons of the first byte and the second byte relative to the plurality of templates. The simultaneous comparison of the first byte with at least the first min term, and the simultaneous comparison of the second byte following the first byte with at least the second min term, ensures that the min terms are sequentially processed based on the order in which the relevant data byte is received. Hence, multiple templates may be processed simultaneously and in real time, by ordering the min terms based on the incoming data stream. Consequently, an incoming data packet may be evaluated at the network switch port in real time, minimizing latency in the network switch port.

Still another aspect of the present invention provides a network switch port filter configured for evaluating an incoming data packet. The network switch port filter includes a min term memory configured for storing min term values. Each min term value is stored in the min term memory based on a corresponding selected byte of the incoming data packet for comparison. The min term value also has an expression portion specifying a corresponding comparison operation, and a template identifier field that specifies templates that use the corresponding min term. The network switch port

filter also includes a min term generator configured for simultaneously comparing a received byte of the incoming data packet with the min terms that correspond to the received byte and generating respective min term comparison results, and an equation core configured for generating a frame tag identifying the incoming data packet based on the min term comparison results relative to the templates. Since each min term value is stored in a location based on the corresponding selected byte of the incoming data packet, each min term can be accessed from the min term memory, and compared with the selected byte, based on the order in which the selected byte is received by the network switch port. Hence, the min term generator can simultaneously compare a given received byte of the incoming data packet with all the relevant min terms, enabling real-time evaluation of the incoming data packet.

Additional advantages and novel features of the invention will be set forth in part in the description which follows and in part will become apparent to those skilled in the art upon examination of the following or may be learned by practice of the invention. The advantages of the present invention may be realized and attained by means of instrumentalities and combinations particularly pointed in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

Reference is made to the attached drawings, wherein elements having the same reference numeral designations represent like element elements throughout and wherein:

Figure 1 is a block diagram of a packet switched network including multiple network switches for switching data packets between respective subnetworks according to an embodiment of the present invention.

Figure 2 is a diagram illustrating a conventional layer 2 Ethernet-type data packet carrying a layer 3 Internet protocol (IP) packet.

Figure 3 is a flow diagram illustrating a conventional (prior art) method of evaluating an IP packet.

Figure 4 is a block diagram illustrating the network switch port of Figure 1, including the network switch port filter, according to an embodiment of the present invention.

Figures 5A and 5B are diagrams illustrating simultaneous processing of four templates of an equation by the min term generator of Figure 4.

Figure 6 is a diagram illustrating in further detail the simultaneous processing of min terms by the min term generator of Figure 4.

Figure 7 is a diagram illustrating a table entry of a min term in the min term memory of Figure 4 according to an embodiment of the present invention.

BEST MODE FOR CARRYING OUT THE INVENTION

Figure 1 is a block diagram illustrating a packet switched network 10, such as an Ethernet (IEEE 802.3) network. The packet switched network includes integrated (i.e., single chip) multiport switches 12 that enable communication of data packets between network stations 14. Each network station 14, 5 for example a client workstation, is typically configured for sending and receiving data packets at 10 Mbps or 100 Mbps according to IEEE 802.3 protocol. Each of the integrated multiport switches 12 are interconnected by gigabit Ethernet links 16, enabling transfer of data packets between subnetworks 18a, 18b, and 18c. Hence, each subnetwork includes a switch 12, and an associated group of network stations 14.

10 Each switch 12 includes a switch port 20 that includes a media access control (MAC) module 22 and a packet classifier module 24. The MAC module 20 transmits and receives data packets to the associated network stations 14 across 10/100 Mbps physical layer (PHY) transceivers (not shown) according to IEEE 802.3u protocol. Each switch 12 also includes a switch fabric 25 configured for making frame forwarding decisions for received data packets. In particular, the switch fabric 25 is 15 configured for layer 2 switching decisions based on source address, destination address, and VLAN information within the Ethernet (IEEE 802.3) header; the switch fabric 25 is also configured for selective layer 3 switching decisions based on evaluation of an IP data packet within the Ethernet packet.

As shown in Figure 1, each switch 12 has an associated host CPU 26 and a buffer memory 28, for example an SRAM. The host CPU 26 controls the overall operations of the corresponding switch 20 12, including programming of the switch fabric 25. The buffer memory 28 is used by the corresponding switch 12 to store data frames while the switch fabric 25 is processing forwarding decisions for the received data packets.

As described above, the switch fabric 25 is configured for performing layer 2 switching decisions and layer 3 switching decisions. The availability of layer 3 switching decisions may be 25 particularly effective if an end station 14 within subnetwork 18a wishes to send an e-mail message to selected network stations in subnetwork 18b, 18c, or both; if only layer 2 switching decisions were available, then the switch fabric 25 of switch 12a would send the e-mail message to switches 12b and 12c without specific destination address information, causing switches 12b and 12c to flood all their ports. Otherwise, the switch fabric 25 of switch 12a would need to send the e-mail message to a router (not 30 shown), which would introduce additional delay. Use of layer 3 switching decisions by the switch fabric 25 enables the switch fabric 25 to make intelligent decisions as far as how to handle a packet, including advanced forwarding decisions, and whether a packet should be considered a high-priority packet for latency-sensitive applications, such as video or voice. Use of layer 3 switching decisions by the switch fabric 25 also enables the host CPU 26 of switch 12a to remotely program another switch, for example 35 switch 12b, by sending a message having an IP address corresponding to the IP address of the switch

12b; the switch 12b, in response to detecting a message addressed to the switch 12b, can forward the message to the corresponding host CPU 26 for programming of the switch 12b.

Figure 2 is a diagram illustrating an Ethernet (IEEE 802.3) packet 30 carrying an IP packet 32 as payload data. Specifically, the Ethernet packet 30 includes a start frame delimiter (SFD) 34, an Ethernet header 36, the IP packet 32, and a cyclic redundancy check (CRC) or frame check sequence (FCS) field 38. Hence, a switch fabric 25 configured for layer 3 switching decisions needs to be able to quickly process the IP packet 32 within the received Ethernet frame 30 to avoid blocking of the frame within the switch.

Figure 3 is a flow diagram illustrating an example of the type of layer 3 processing that might be performed for an incoming data packet. The flow diagram of Figure 3, conventionally implemented in software, would involve checking whether the incoming data packet was a hypertext transport protocol (HTTP) packet in step 50, an SNMP packet in step 52, or a high-priority packet in step 54. The appropriate tag would then be assigned identifying the packet in steps 56, 57, 58, or 60.

The arrangement of Figure 3, however, cannot from a practical standpoint be implemented in hardware in a manner that would provide a non-blocking switch for 100 Mbps or gigabit networks. In particular, the sequential nature of the decision process in Figure 3 would result in undue latency for the incoming data packet.

According to the disclosed embodiment, the packet classifier module 24 of Figure 1 is configured for multiple simultaneous comparisons between the incoming data stream and templates that identify the data format of the incoming data stream. Specifically, users of the host processor 26 will specify policies that define how data packets having certain IP protocols should be handled by the switch fabric 25. These policies are implemented by loading into the switch fabric 25 a set of frame forwarding decisions for each corresponding IP protocol type. Hence, the switch fabric 25 could include one set of frame forwarding instructions for an HTTP packet, another set of frame forwarding instructions for an SNMP packet, and another set of frame forwarding instructions for a high-priority packet (e.g., video, or voice, etc.).

Figure 4 is a block diagram illustrating the packet classifier module 24 according to an embodiment of the present invention. As shown in Figure 4, the network switch port 20 includes a MAC 22, a receive FIFO buffer 27, a header modifier 29, and the packet classifier module 24. The packet classifier module 24, also referred to as a network switch port filter, is configured for identifying (i.e., evaluating) the incoming data packet at the network switch port 20, and supplying to the switch fabric 25 a tag that specifies the action to be performed on the data packet based on type of data packet being received. Specifically, the packet classifier module 24 simultaneously compares the incoming data packet with a plurality of templates configured for identifying respective data formats. The packet classifier module 24, based on the comparison between the incoming data packet and the plurality of

templates, identifies an equation to be executed that specifies the tag to be supplied to the switch fabric 25.

Specifically, the packet classifier module 24 generates a comparison result that identifies the incoming data packet by detecting at least one matched template from a plurality of templates. The 5 packet classifier module 24 then identifies which of the equations includes the matched template, and generates the tag specified by the equation.

Figures 5A and 5B are diagrams illustrating the simultaneous processing of two templates of an equation by the packet classifier module 24. Figure 5A illustrates the logical evaluation by the packet classifier module 24 of the equation:

10 $Eq1 = M1 * M2 * M3 * M4 * (M5 + M6 + M7 + M8)$.

Figure 5B illustrates how the equation Eq1 would actually be stored in the min term memory 70. The equation Eq1 includes four templates 62a, 62b, 62c, and 62d: the template 62a includes the min terms M1, M2, M3, M4, and M5; the template 62b includes the min terms M1, M2, M3, M4, and M6; 15 the template 62c includes the min terms M1, M2, M3, M4, and M7; and the template 62d includes the min terms M1, M2, M3, M4, and M8. Each template 62 corresponds to a specific IP data format recognizable based on the header of the IP data packet 32. For example, templates 62a and 62c may be configured for identifying an HTTP packet, and templates 62b and 62d be may be configured for identifying an SNMP packet. Specifically, an HTTP packet is identified if it is in IPv4 format, the time 20 to live field in IP is bigger than one, the protocol field in IP header is TCP, header checksum is correct, source TCP port is 80 or destination TCP port is 80. An SNMP packet is identified if it is in IPv4 format, the time to live field in IP is bigger than one, the protocol field in IP header is TCP, header checksum is correct, source TCP port is 25 or destination TCP port is 25.

Hence, the following min terms may be established to represent all the above-described criteria:

- 25 M1= packet is in IPv4 format
 M2= time to live field in IP is bigger than one
 M3= protocol field in IP header is TCP
 M4= header checksum is correct
 M5= source TCP port is 80
30 M6= destination TCP port is 80
 M7= source TCP port is 25
 M8= destination TCP port is 25

Hence, the templates 62a and 62c identify HTTP packets, and the templates 62b and 62d identify SNMP packets. Thus, equation one (Eq1) specifies that a specific result (e.g., the tag having a specified value) should be output to the switch fabric 25 if either template 62a, 62b, 62c, or 62d are true.

Moreover, the min terms M1...M8 are arranged within the associated templates 62a and/or 62b

- 5 in a prescribed order that corresponds to the relative position of a data byte in the incoming data stream. As illustrated in Figure 6, the min term M1 is configured for comparison with the first byte (B1) of the IP packet 32, the min term M2 is configured for comparison with a subsequent byte (B2) of the IP packet 32 that follows B1, the min term M3 is configured for comparison with a subsequent byte (B3) that follows B2, etc. Hence, the use of templates 62 having min terms in an order based on the relative position of a
10 data byte in the incoming data stream enables multiple simultaneous comparisons between the incoming data stream and min terms. Hence, an incoming data packet can be compared to multiple templates to determine not only the data format of the incoming data packet, but also what action needs to be performed by the switch fabric 25.

Figure 4 is a block diagram illustrating in detail the packet classifier 24 of Figure 1. As shown

- 15 in Figure 4, the packet classifier 24, also referred to as a network switch port filter, includes a min term memory 70 for storing the min term values (e.g., M1, M2, etc.) as illustrated in Figure 7, described below. The packet classifier 24 also includes a frame identifier 72 configured for identifying the type of layer 2 frame being received; in particular, identifying the type of layer 2 frame being received (e.g., Ethernet, IEEE 802 to 3, etc.) enables identification of the start position 64 of the
20 IP packet 32 within the layer 2 packet 30. The packet classifier 24 also includes a min term controller 74, a min term generator 76, an equation core 78, and an evaluation results memory 80. A processor interface module (pi_mod) 82 is used for transferring the generated min terms from the host CPU 26 into the min term memory 70.

The min term controller 74 is configured for fetching the min terms from the min term

- 25 memory 70 corresponding to a selected byte of the IP frame 32. The min term controller 74 also includes a location converter configured for specifying the actual byte location (byte_location) of the start point 64 in response to receiving a frame type (frm_type) signal from the frame identifier 72 that specifies the type of layer 2 frame. Hence, the min term controller 74, in response to detecting the beginning of the IP packet, fetches all the min terms that are to be compared with the first byte (B1) of
30 the IP packet 32, for example min terms M1, M9, and M14 for equations Eq1, Eq2, and Eq3 in Figure 6. The min term controller 74 then forwards the min term values (M_STRU_INFO) to the min term generator 76 and the equation core 78.

The min term generator 76 performs the actual min term comparisons between the min terms

fetched by the min term controller and the selected byte of the incoming data stream. For example,

- 35 the min term generator 76 simultaneously compares in Figure 6 the incoming data byte B1 with the

min terms M1, M9, and M14, and provides the min term comparison results (mt_result) to the equation core 78. During the next comparison cycle, the min term generator 76 simultaneously compares the incoming data byte B2 with the min terms M2, M10, and M15. According to the disclosed embodiment, the min term generator is configured for simultaneously comparing the 5 incoming data stream to up to eight min terms.

The equation core 78 is configured for generating a frame tag based on the min term comparison results received from the min term generator 76, relative to the relevant templates 62. For example, the equation core 78 evaluates equation 1, illustrated in Figure 5, by evaluating the min term results sequentially as the results are supplied from the min term generator. For example, if the 10 comparisons for each of the min terms M1, M2, M3, M4, M5, and M6 result in a true condition, described below with respect to Figure 7, then the end condition is matched in equation 1, causing the equation core 78 to generate a tag corresponding to the condition specified for equation 1. The frame tag identifies the nature of the incoming data packet, as well as the action that needs to be performed by the switch fabric 25.

15 Figure 7 is a diagram illustrating the data format of the min term structure in the min term memory 70. According to the disclosed embodiment, the min terms are stored in the min term memory 70 in an order corresponding to the relative position of the data byte to be compared. Hence, all min terms that are to be compared to the first data byte are stored together in a first part of the min term memory, followed by min terms to be compared with the second data byte, etc.

20 Alternatively, the min terms may be stored in an order based on relevant information within the IP header, as well as the relative position of the data byte to be compared. Hence, the min terms may be stored in an order for comparing the sequence of data bytes providing the source IP address, destination IP address, and source and source and destination ports; in this case, non-relevant data bytes at the beginning of the IP frame would not have associated min terms stored at the beginning of 25 the min term memory 70, further improving the efficiency of the min term memory 70.

Each table entry 90 includes a min term portion and an evaluation portion. The min term portion includes a mask field (MASK) 94, an expected data field (EXP_DATA) 96, and an operator field (OPERATOR) 98. Based on the position of the table entry 90 in the min term memory 70, the min term controller 74 is able to determine which byte of the IP packet 32 that needs to be compared 30 with the corresponding min term, relative to the beginning 64 of the IP packet. The mask field 94 is a mask that is used by the min term generator 76 in performing comparisons; if the mask has a bit set to 1, the value is compared, and if the mask value has zeros in the field, the comparison is a don't care. The expected data field 96 specifies the expected data to be compared with the relevant data byte of the IP packet 32. The operator field 98 specifies the type of comparison to be performed by the min

term generator, for example: less than, less than or equal to, equal to, greater than, greater than or equal to, and not equal to.

The evaluation portion includes a branches portion 100, a response portion (RINP1) 102 for the case where the comparison of the min term portion is true, a second response portion (RINP0) 106 for the case where the comparison of the min term portion is false, and an equation identifier 110. 5 The branches portion 100 specifies the order of the OR term in the equation; for example, the min term M1 as shown in Figures 5 and 6 would have its branches portion set to 0000 1111, indicating that the first four branches of the equation specified in the equation identifier field 110 are to include the corresponding min term. The use of eight bits for the branches portion assumes that there are a 10 maximum of eight branches in any given equation.

The response portion 102 specifies the operation to be performed if the min term portion is evaluated as true relative to the compared data byte. In particular, the finish bit (FIN) is set to one if the results of the equation is determined if the min term result is true; the back to initial (BINIT) is set to one if the evaluation process should return to the initial state (init) if the min term result is true. For 15 example, in the case of min term M1, the FIN bit and the BINIT bit of RINP1 are set to zero, since additional comparisons are needed if the min term result is true. In the case of min terms M5, M6, M7, and M8, the FIN bit of RINP1 is set to one, since a comparison result of "true" results in the end of the evaluation, as shown in Figure 5.

The response portion 106 specifies the operation to be performed if the min term portion is 20 evaluated as false relative to the compared data byte. In particular, the finish bit (FIN) is set to one if the results of the equation is determined if the min term result is false; the back to initial (BINIT) is set to one if the evaluation process should return to the initial state (init) if the min term result is false. For example, in the case of min term M1, the FIN bit is set to zero and the BINIT bit of RINP1 is set to one, such that the equation would return to the INIT state if the min term result M1 was false, as 25 shown in Figure 5.

The equation identifier field 110 identifies the equation (or template if there is only one template in an equation) that the min term corresponds to.

Hence, the equation core 78 determines whether any specified equation has a template 62 that matches the incoming data stream. Based on the multiple simultaneous comparisons of the incoming 30 data stream with the multiple templates 62, the equation core 78 can identify a matching equation, and generate the appropriate tag corresponding to the matched equation for help it to the switching fabric 25. If desired, the core 78 by also output a command to the header modifier 29 to modify the layer 2 header, the layer 3 header, or both, before transferring the data to the switch.

According to the disclosed embodiment, a network switch port includes a filter capable of 35 performing multiple simultaneous comparisons between the incoming data stream of the data packet

and multiple templates configured for identifying a corresponding protocol. Since the packet classifier module 24 can process any of the bytes of the IP frame 32, the packet classifier module 24 can interpret all the header information in the IP packet 32 from layer 3 up to layers 7 protocols. Moreover, the multiple simultaneous comparisons enables the network switch 12 to perform layer 3
5 switching for 100 Mbps and gigabit networks without blocking in the network switch. Finally, the multiple simultaneous comparisons in the order in which the data is received enables real time comparisons to be performed, as opposed to alternative schemes such as programmable logic arrays (PLAs), which would require the entire header to be received before processing can begin.

While this invention has been described with what is presently considered to be the most
10 practical preferred embodiment, it is to be understood that the invention is not limited to the disclosed embodiments, but, on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims.

What is Claimed Is:

1. A method of evaluating an incoming data packet at a network switch port, the method comprising:

storing a plurality of templates configured for identifying respective data formats, each template having at least one min term configured for comparing a corresponding prescribed value to a corresponding selected byte of the incoming data packet;

simultaneously comparing, to the selected byte, the min terms that correspond to the selected byte as the selected byte is received by the network switch port; and

generating a comparison result that identifies the incoming data packet, based on the comparisons of the min terms to the data bytes received by the network switch port.

2. The method of claim 1, wherein the simultaneously comparing step includes:

loading the min terms corresponding to a first of the data bytes into a min term generator;

comparing in parallel the min terms loaded in the min term generator with the first of the data bytes; and

outputting comparison results for the min terms loaded in the min term generator to an evaluation core.

3. The method of claim 2, wherein the simultaneously comparing step further includes loading the min terms corresponding to a second of the data bytes, contiguously following the first of the data bytes, into the min term generator.

4. The method of claim 1, further comprising outputting a frame tag, based on the comparison result, to a switch fabric configured for selectively switching the incoming data packet based on the corresponding frame tag.

5. The method of claim 1, wherein the storing step includes storing each min term in a memory as a table entry, each table entry having a location in the memory based on a location of the corresponding selected byte in the incoming data packet, the table entry including a min term expression portion specifying the corresponding prescribed value and a comparison operand, and a template identifier field that specifies the templates that use the corresponding min term.

6. The method of claim 5, wherein the generating step includes:

temporarily storing results of the comparisons of the min terms to the selected bytes of the incoming data packet;

detecting at least one matched template from the plurality of templates based on the results of the comparisons of the min terms; and

generating the comparison result based on the detected at least one matched template.

7. The method of claim 5, wherein the first of the data bytes corresponds to a first of the data bytes of a packet having a prescribed format, the simultaneously comparing step including evaluating the selected data byte relative to a beginning of the packet having the prescribed format.

8. The method of claim 7, wherein the prescribed format is Internet protocol (IP) format.

9. The method of claim 6, wherein the step of generating the comparison result based on the detected at least one matched template includes:

identifying for each of the min terms compared with the incoming data packet a corresponding equation, each equation specifying a unique result for a selected group of the templates; and

generating the comparison result by the equation having the detected at least one matched template.

10. A method of evaluating an incoming data packet at a network switch port, the method comprising:

simultaneously comparing a first byte of the incoming data packet as the first byte is received by the network switch port, with a plurality of templates by comparing the first byte with at least a first min term associated with at least one of the templates, each template configured for identifying a prescribed format within the incoming data packet;

simultaneously comparing a second byte of the incoming data packet, following the first byte, with the templates by comparing the second byte with at least a second min term associated with at least one of the templates; and

generating a comparison result based on min term results from the comparisons of the first byte and the second byte relative to the plurality of templates.

11. The method of claim 10, wherein the first min term is associated with a first group of the templates and a third min term, configured for comparison with the first byte, is associated with a second group of the templates, the step of simultaneously comparing the first byte including simultaneously comparing the first byte with the first and third min terms.

12. The method of claim 11, wherein each template has a plurality of min terms for comparison with respective bytes of the incoming data packet, the method further comprising successively comparing the min terms of the templates based on the relative positions of the respective bytes of the incoming data packet.

13. The method of claim 12, wherein the generating step includes generating an equation result specifying an action to be performed by a network switch fabric, based on the comparison results for the templates assigned to the equation result.

14. A network switch port filter configured for evaluating an incoming data packet, comprising:

a min term memory configured for storing min term values, each min term value stored based on a location of a corresponding selected byte of the incoming data packet for comparison, an expression portion specifying a corresponding comparison operation, and a template identifier field that specifies templates that use the corresponding min term;

a min term generator configured for simultaneously comparing a received byte of the incoming data packet with the min terms that correspond to the received byte and generating respective min term comparison results; and

an equation core configured for generating a frame tag identifying the incoming data packet based on the min term comparison results relative to the templates.

15. The filter of claim 14, further comprising a frame identifier configured for identifying a type of layer 2 packet, the selected byte of the incoming data packet determined based on the identified type of layer 2 packet.

16. The filter of claim 15, wherein the location of each stored min term value is relative to a beginning of an IP frame within the layer 2 packet.

17. The filter of claim 16, further comprising a min term controller configured for fetching the min terms from the min term memory corresponding to a selected byte of the IP frame within the incoming data packet.

18. The filter of claim 15, wherein the equation core generates the frame tag at a wire rate of the incoming data packet and prior to an end of the incoming data packet.

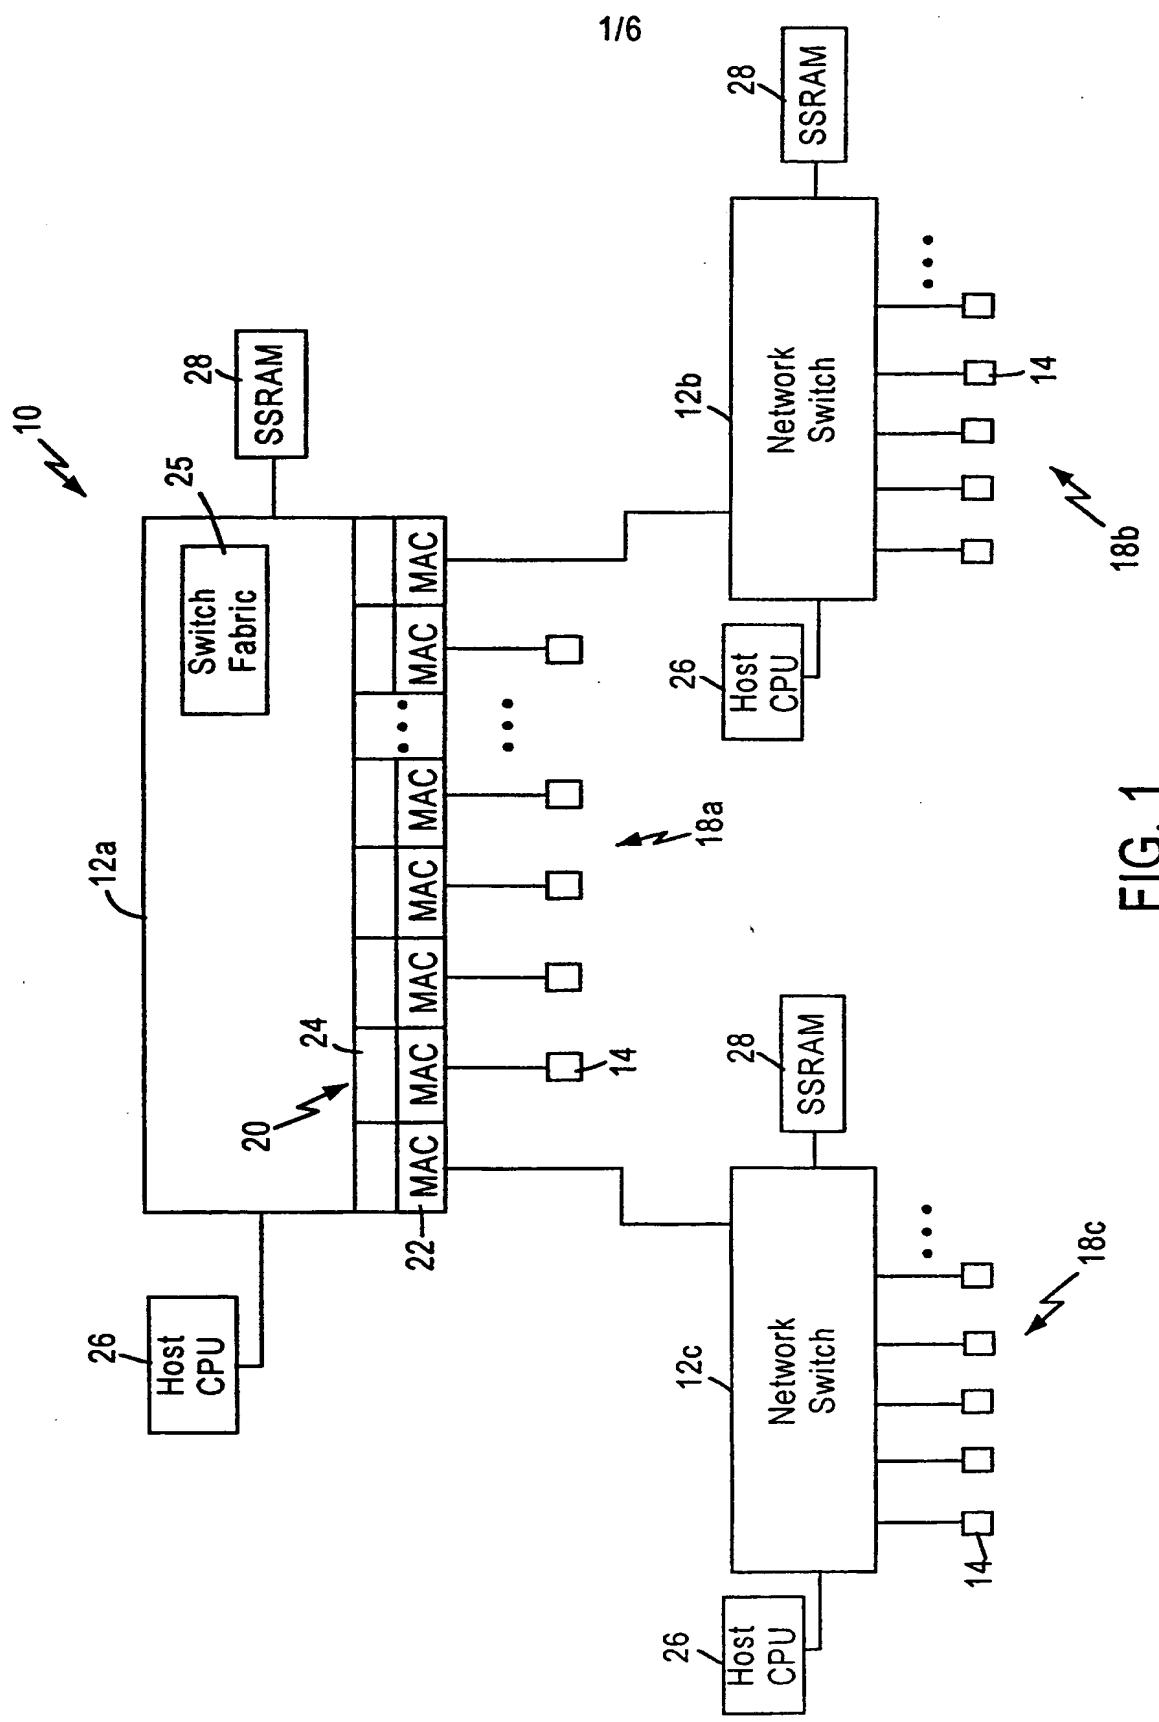


FIG. 1

2/6

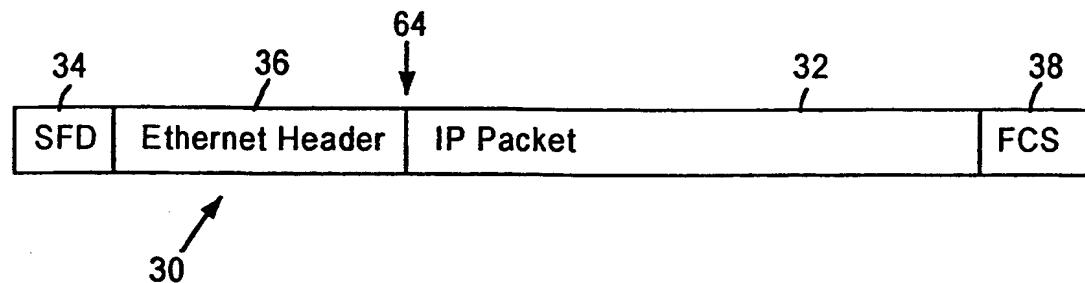


FIG. 2

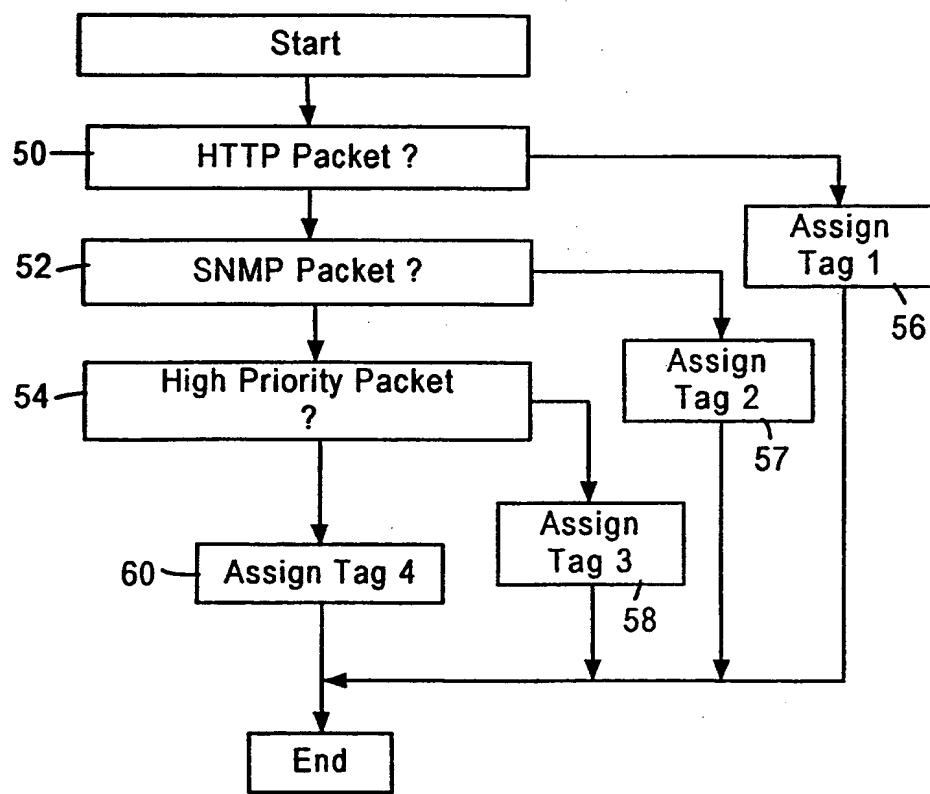


FIG. 3

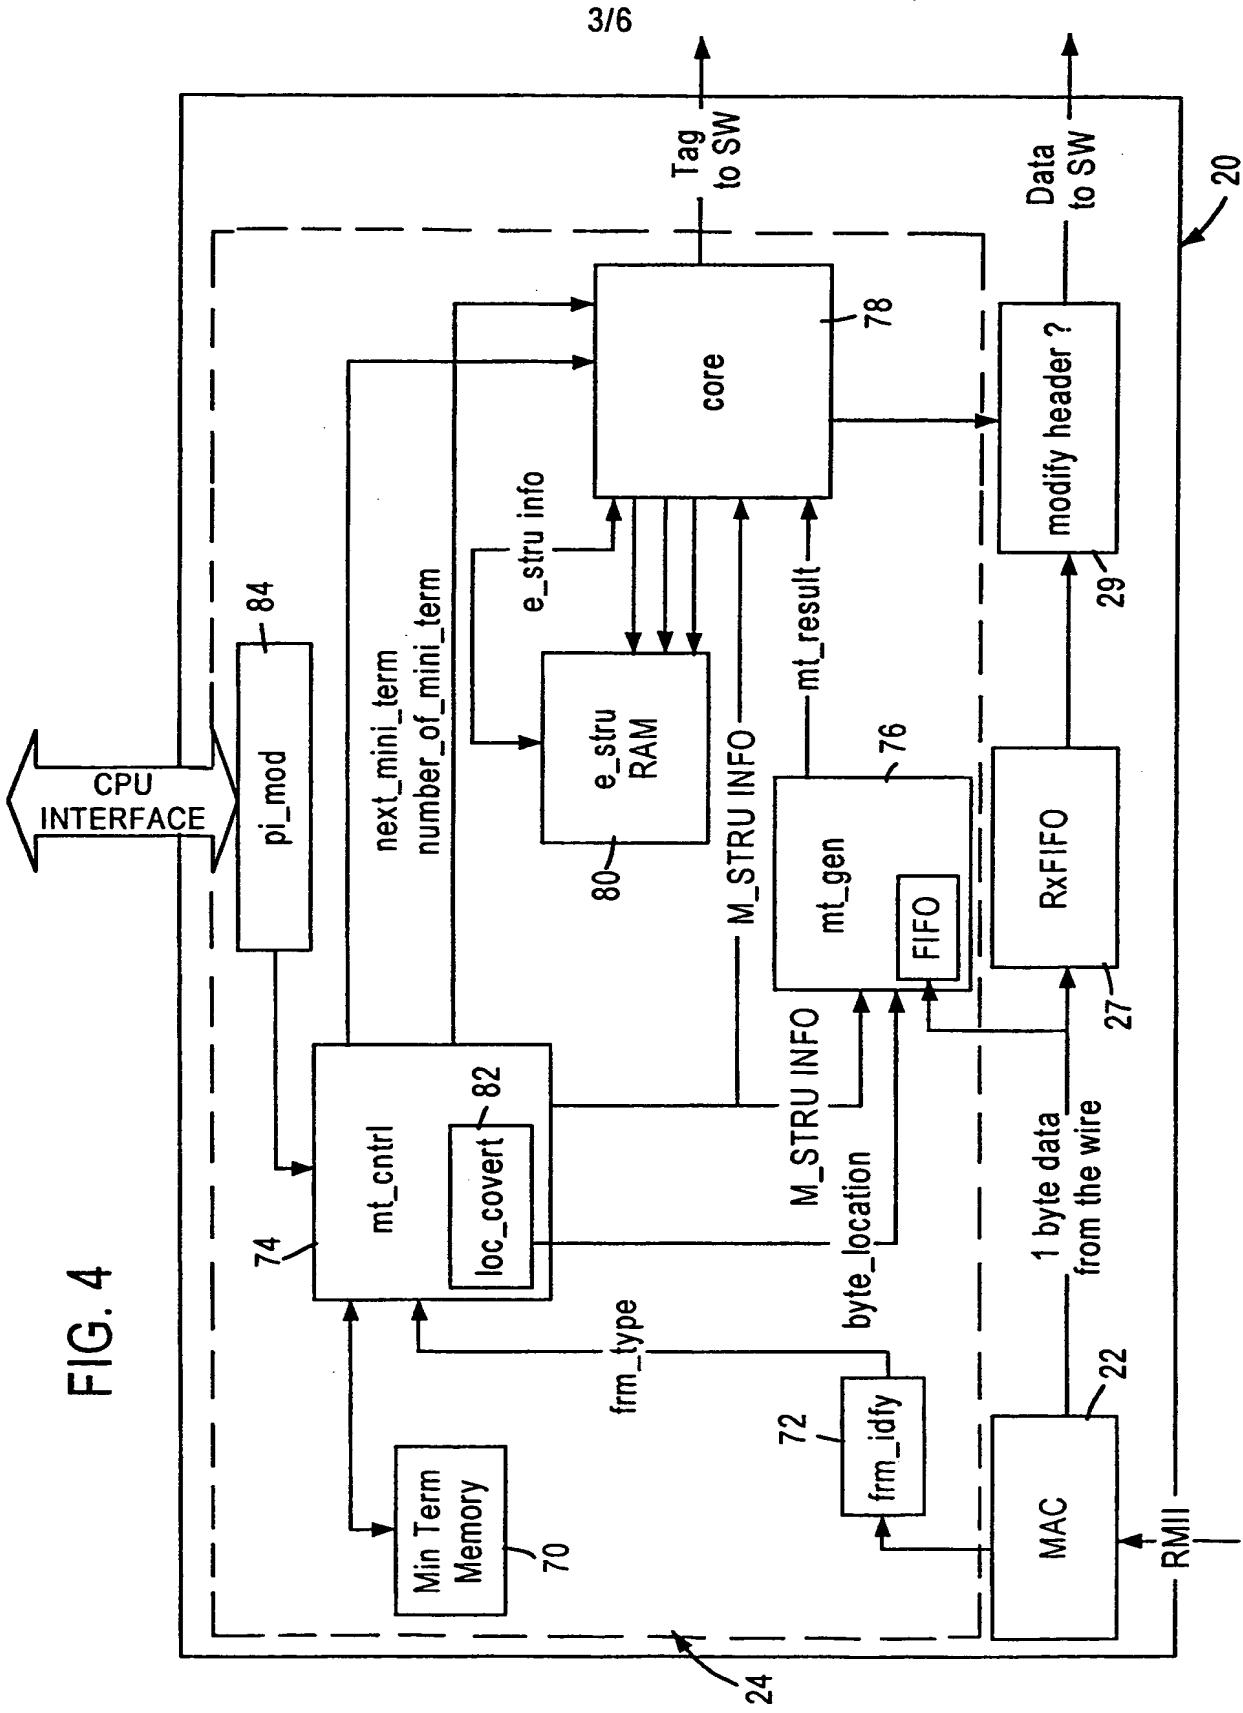
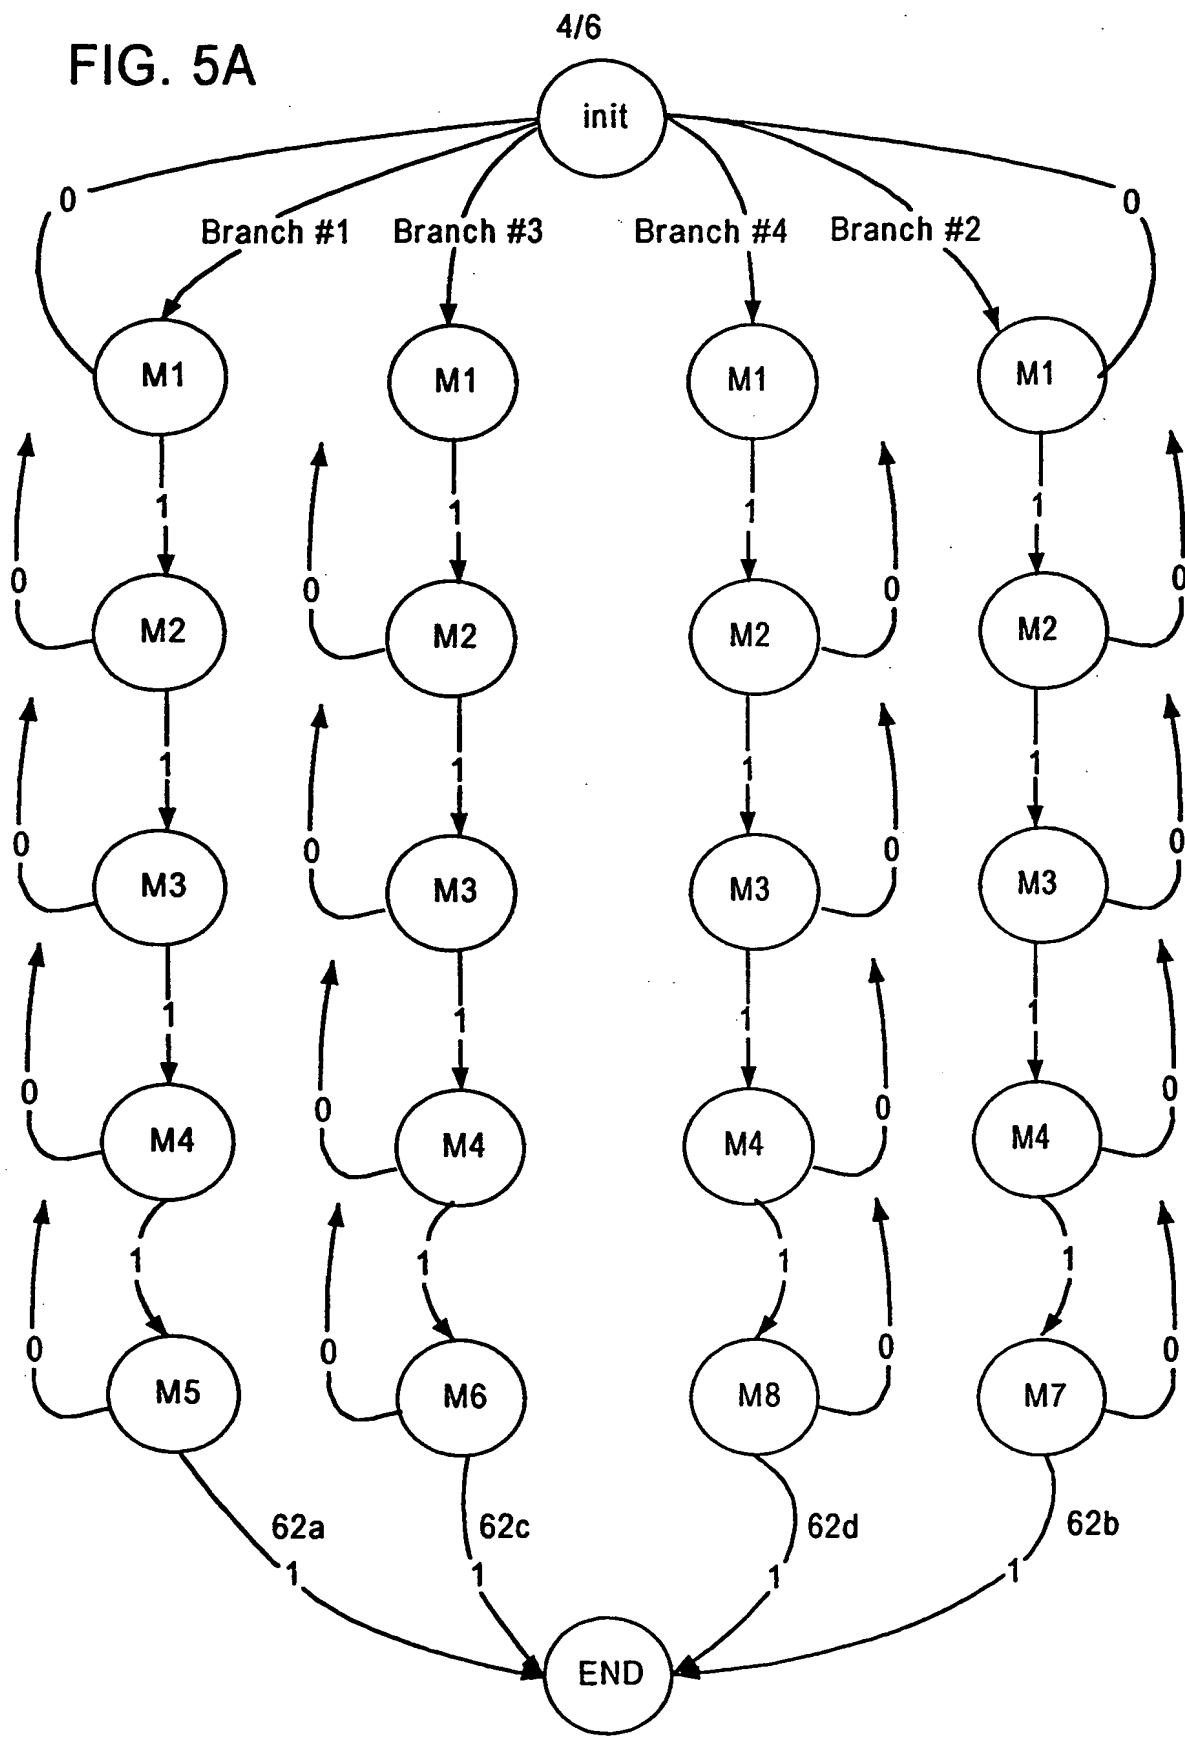
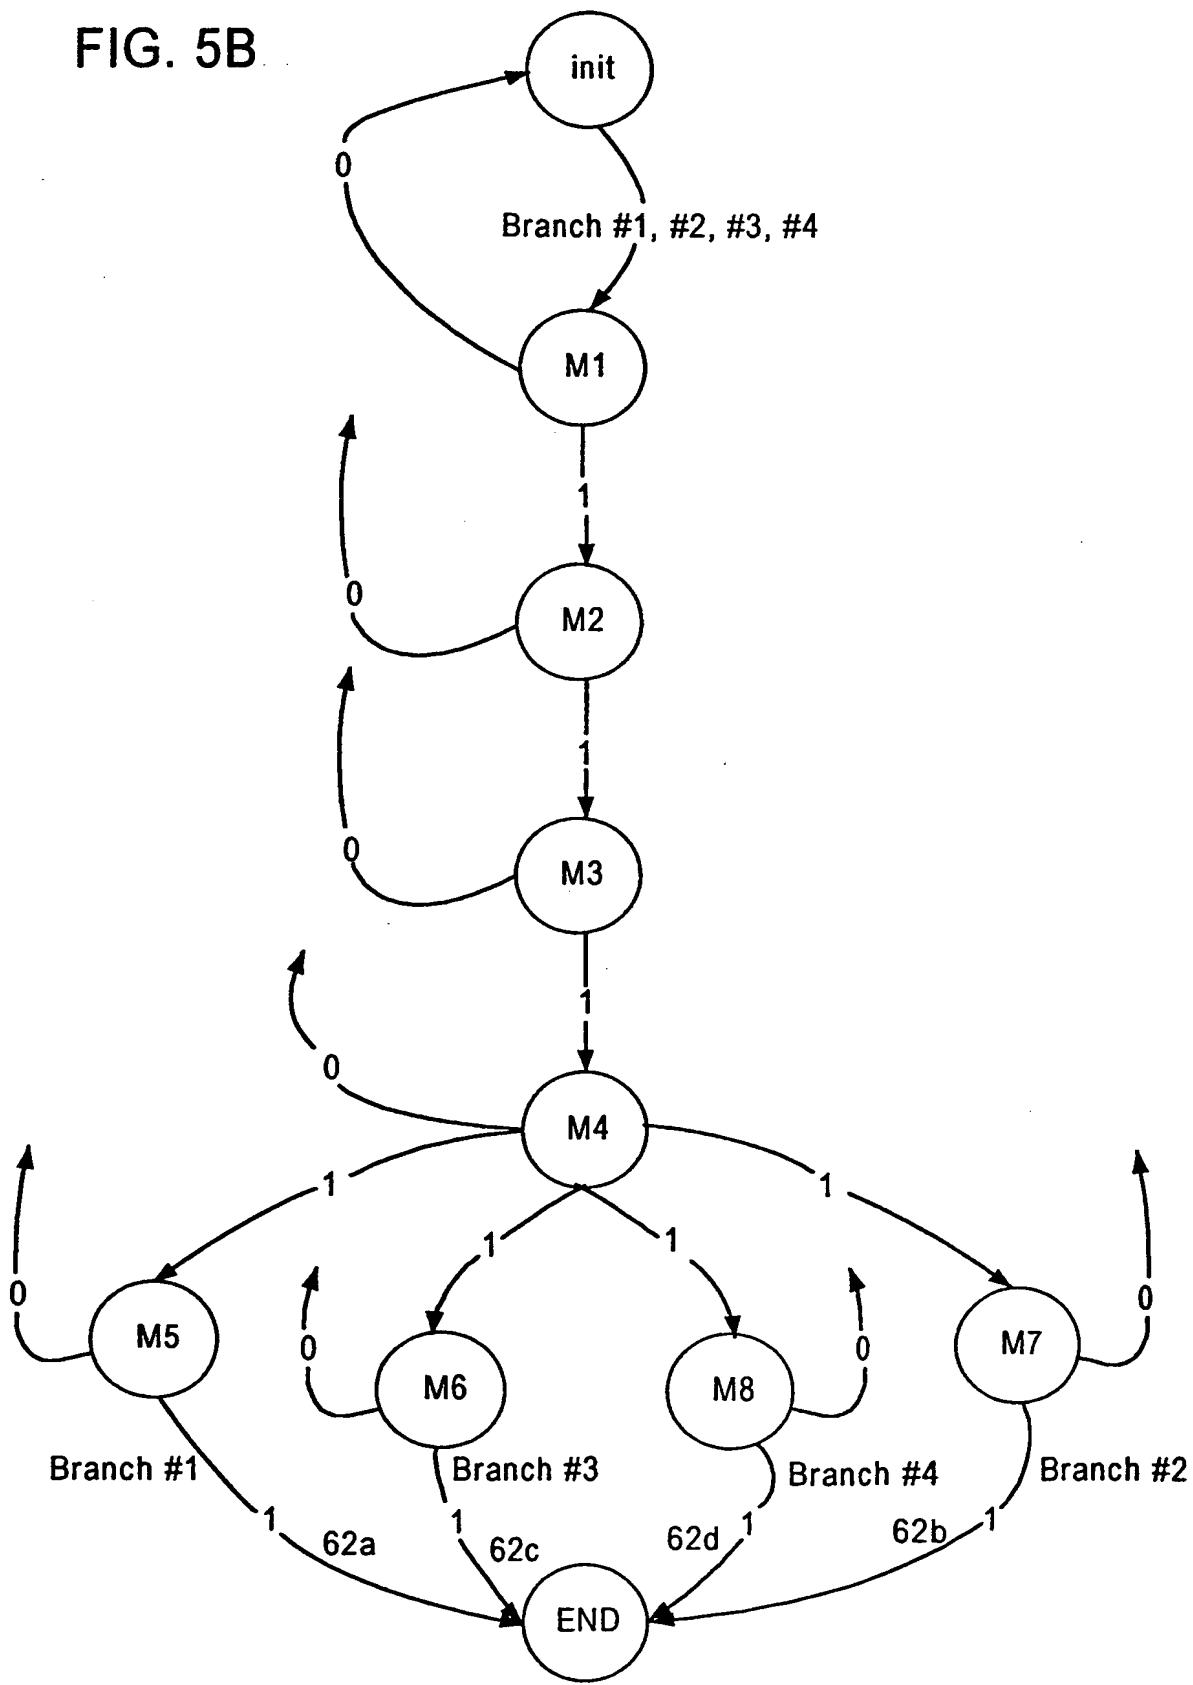


FIG. 5A



5/6

FIG. 5B



6/6

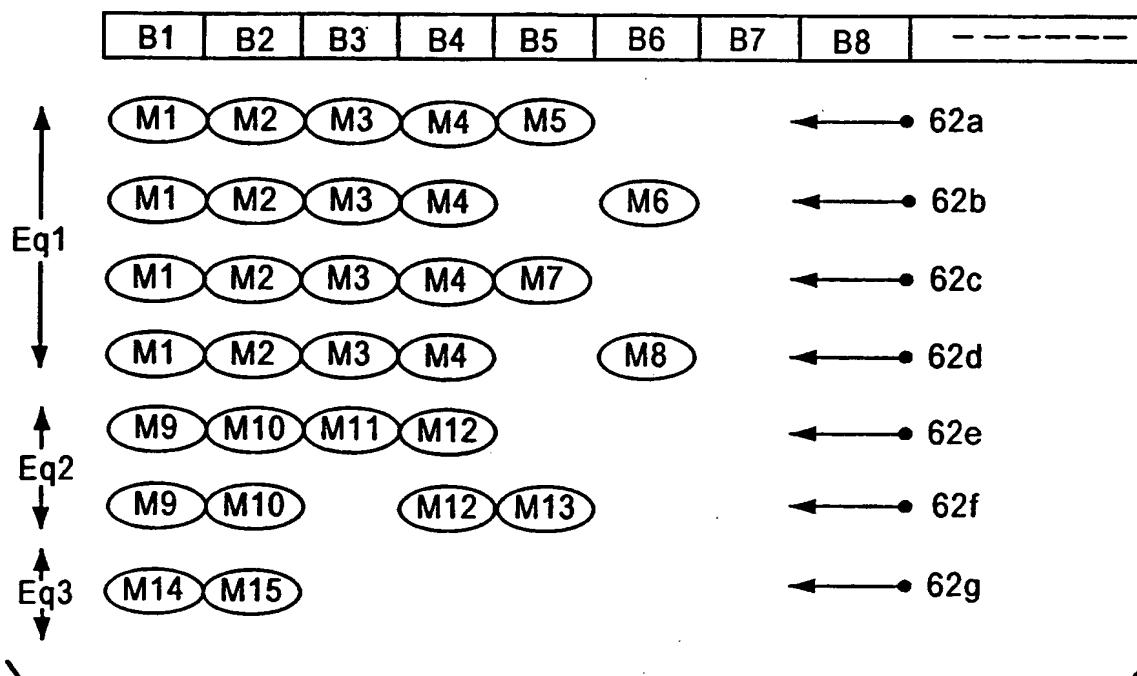


FIG. 6

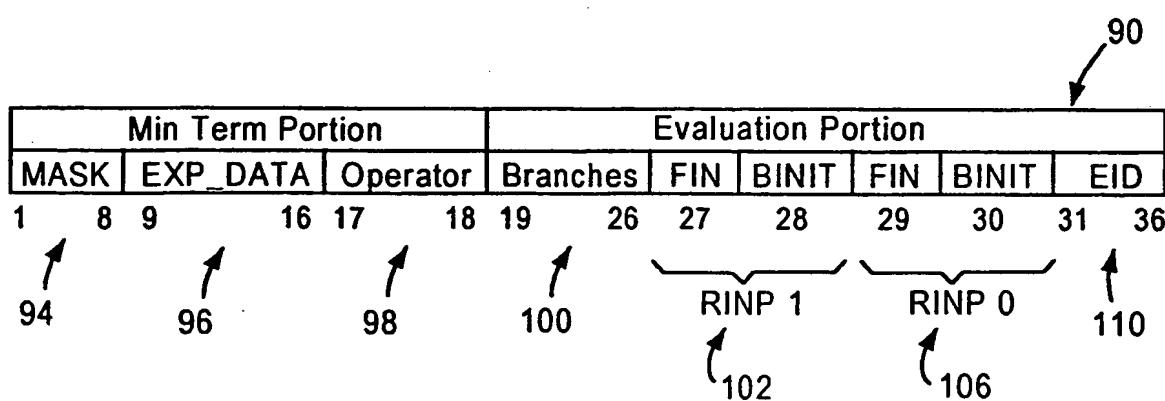


FIG. 7

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/13397

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04L12/46 H04L12/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>WO 98 35480 A (BARES WILLIAM H ;DEB ALAK K (US); XAQTI CORP (US); SAMBAMURTHY NAM) 13 August 1998 (1998-08-13)</p> <p>abstract</p> <p>page 7, line 10 - line 19 page 10, line 1 -page 11, line 10 page 34, line 14 - line 21 page 20, line 20 -page 21, line 25 page 22, line 24 -page 23, line 12 page 24, line 22 -page 26, line 14</p> <p>-----</p> <p>-/-</p>	1-18

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

4 December 2000

08/12/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
 Fax: (+31-70) 340-3016

Authorized officer

Perez Perez, J

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/13397

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 99 53648 A (TOP LAYER NETWORKS INC) 21 October 1999 (1999-10-21) abstract page 2, line 36 -page 4, line 17 page 13, line 25 -page 14, line 15 page 16, line 26 - line 34 page 22, line 17 - line 24 page 29, line 3 - line 9 page 30, line 3 - line 30 ---	1-18
Y	US 5 546 387 A (LARSSON LEIF M ET AL) 13 August 1996 (1996-08-13) column 3, line 24 -column 4, line 27 column 7, line 9 - line 22 ---	1-7, 10-15,18
Y	US 5 802 054 A (BELLINGER DONALD M) 1 September 1998 (1998-09-01) column 3, line 25 - line 39 column 4, line 23 - line 39 column 13, line 62 -column 14, line 55 ---	1-7, 10-15,18
A	"VIRTUAL LANs GET REAL ETHERNET SWITCH MAKERS ARE TAKING THE LEAD IN DEPLOYING VIRTUAL LANs ACROSS CAMPUS NETWORKS" DATA COMMUNICATIONS, US, McGRAW HILL, NEW YORK, vol. 24, no. 3, 1 March 1995 (1995-03-01), pages 87-92,94,96., XP000496027 ISSN: 0363-6399 page 88, middle column, line 1 - line 13 page 90, middle column, line 9 - line 20 page 96, framed section ---	1-18

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/13397

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
WO 9835480 A	13-08-1998	AU 6169398 A AU 6656998 A EP 0960517 A WO 9835307 A US 6108713 A		26-08-1998 26-08-1998 01-12-1999 13-08-1998 22-08-2000
WO 9953648 A	21-10-1999	AU 3638999 A		01-11-1999
US 5546387 A	13-08-1996	AU 696446 B AU 6266094 A BR 9406380 A CA 2157877 A CN 1119059 A EP 0688484 A FI 954217 A JP 8510102 T KR 237336 B NO 953538 A SE 9300793 A WO 9421069 A		10-09-1998 26-09-1994 16-01-1996 15-09-1994 20-03-1996 27-12-1995 08-09-1995 22-10-1996 15-01-2000 09-11-1995 11-09-1994 15-09-1994
US 5802054 A	01-09-1998	US 5949786 A		07-09-1999

THIS PAGE BLANK (USPTO)